IPv6 Multicast

Δημήτριος Α Αδάμος

ΑΠΘ - ΕΔΕΤ

Agenda

- **Why IPv6?**
- **IPv6 multicast addresses**
- **IPv6 multicast features**
- **Configuration**

Why IPv6?

Higher Ed./Research

- Media services
- Collaboration
- Mobility

Consumer

- Set-top Boxes
- Gaming
- Appliances
- Voice/Video
- Security Monitoring

Manufacturing

- Embedded Devices
- Industrial Ethernet
- IP-enabled components

Government (Federal/Public Sector)

- DoD
  - WIN-T
  - FCS
  - JTRS
  - GIG-E
- State/Local

Transportation

- Telematics
- Traffic Control
- Hotspots
- Transit services

Agriculture/Wildlife

- Animal Tags
- Imagery
- Botanical
- Weather

Medical

- Home care
- Imaging
- Mobility

Why IPv6?

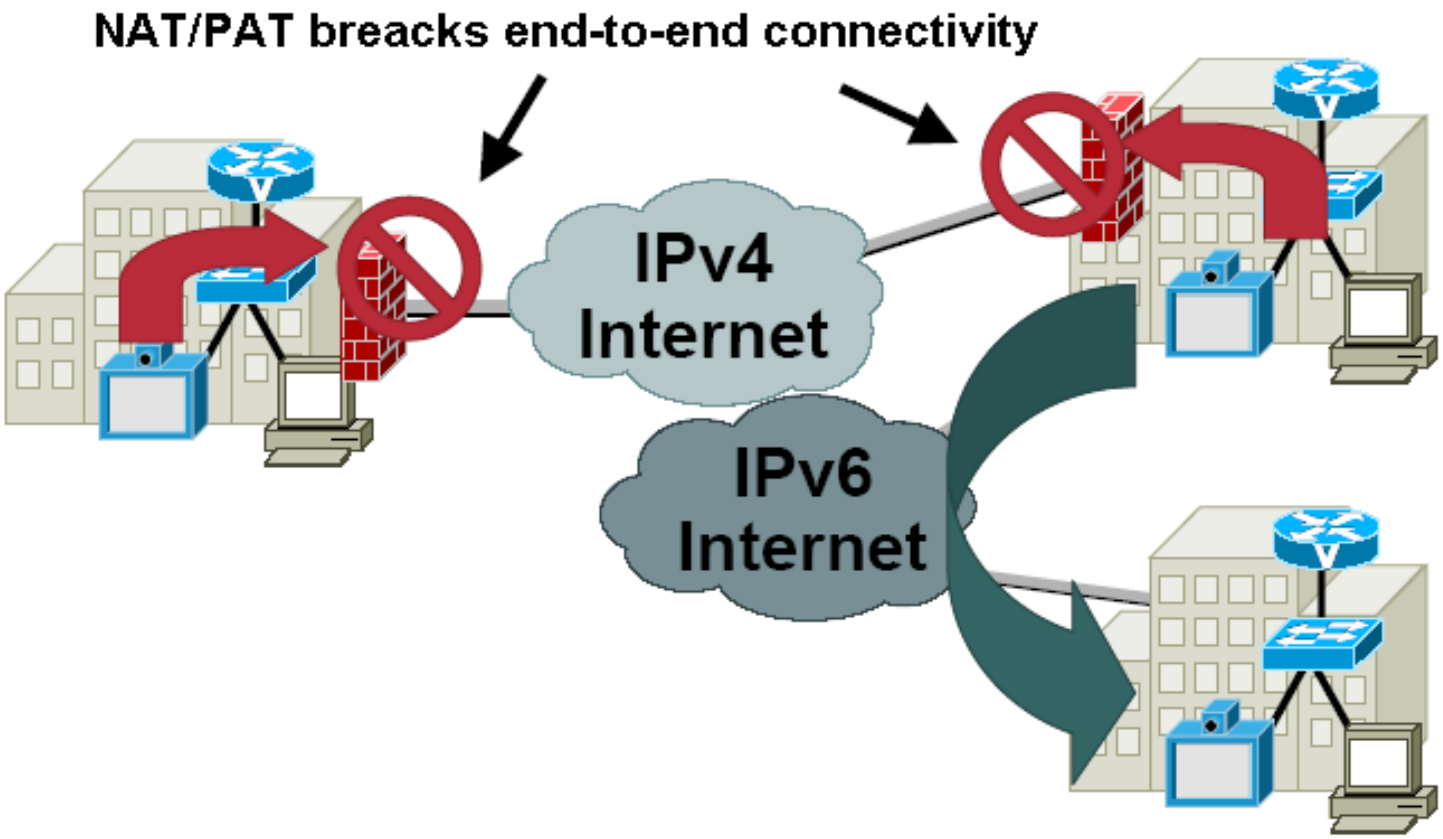

IPv6 restores end-to-end multimedia collaboration

Why IPv6 multicast?

- Efficiently deploy and scale distributed group applications across the network
- Enterprise-wide content distribution model
- Solve traffic congestion problems
- Specific multicast benefits such as scope management

Agenda

- Why IPv6?
- IPv6 multicast addresses
- IPv6 multicast features
- Configuration

IPv6 multicast addresses

- **IPv6 uses a 16bytes/128 bits address length**
- **3 types of address in IPv6:**

  Unicast: one-to-one with various scopes (I.e.:Global,Link,Unique Local)

  Anycast: one-to-nearest (allocated from unicast)

  Multicast: one-to-many

  There is no longer a Broadcast address!

- **All Multicast addresses begin with the format prefix 1111 1111 – easily written as FF**
- **Due to address length, often are sequences that include long zero series. Users often compress zeros as:**

  FF05:0:0:0:0:0:0:2 => FF05::2

IPv6 multicast addresses

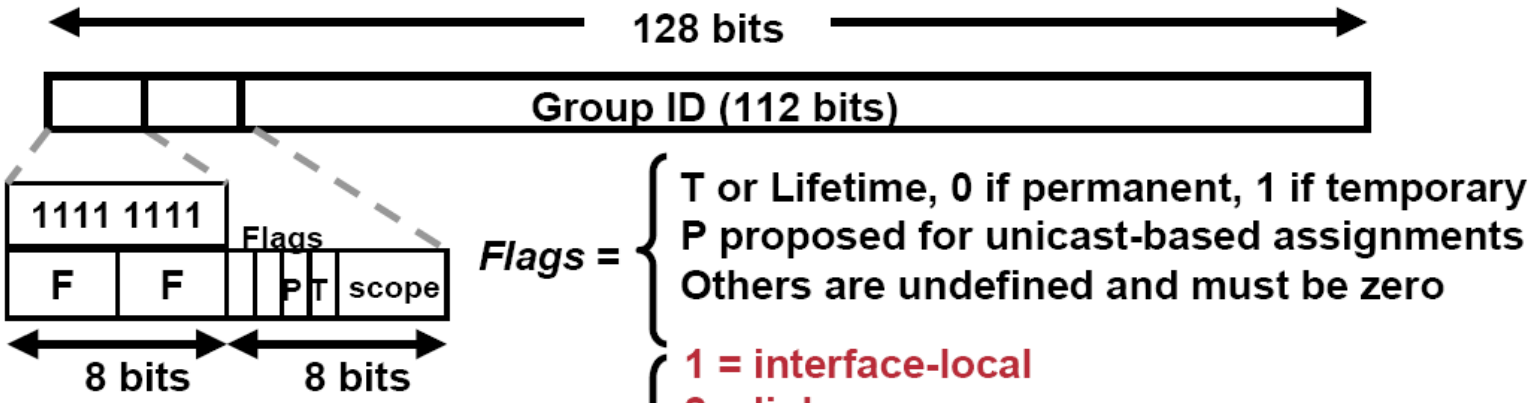

Flags =
- T or Lifetime, 0 if permanent, 1 if temporary
- P proposed for unicast-based assignments
- Others are undefined and must be zero

Scope =
- 1 = interface-local
- 2 = link
- 4 = admin-local
- 5 = site
- 8 = organization
- E = global

Scope in the Network Hierarchy

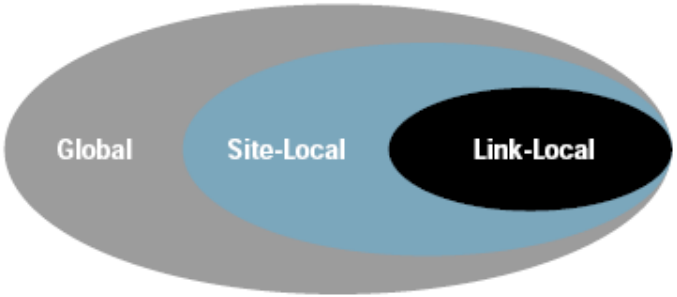

IPv6 actually has possible 15 scopes, from 0–F in hexadecimal, some of which are unused.

For example of an address with the scope set is the 'all' routers addresses below:

FF01:0:0:0:0:0:0:2 with a scope of the local node.

FF02:0:0:0:0:0:0:2 with a scope of the local link.

IPv6 multicast addresses

The "meaning" of a permanently-assigned multicast address is independent of the scope value. For example, if the "NTP servers group" is assigned a permanent multicast address with a group ID of 101 (hex), then:

- FF01:0:0:0:0:0:0:101 means all NTP servers on the same node as the sender.
- FF02:0:0:0:0:0:0:101 means all NTP servers on the same link as the sender.
- FF05:0:0:0:0:0:0:101 means all NTP servers at the same site as the sender.
- FF0E:0:0:0:0:0:0:101 means all NTP servers in the internet

IPv6 Multicast addresses

In IPv6, if you own an IPv6 unicast address prefix you implicitly own an RFC3306 IPv6 multicast address prefix:

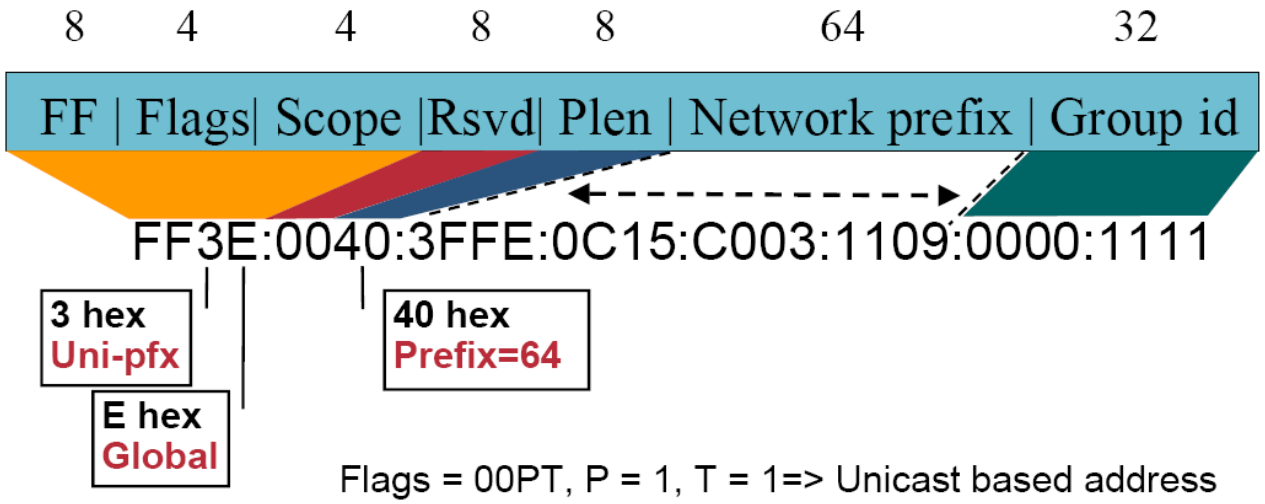

- **RFC 3306 Unicast-Prefix-based IPv6 Multicast**

  The P flag indicates a prefix. Within IPv6 multicast, this flag allows part of the group address to include the source network's Unicast prefix, which creates a globally unique Group Address.

- **Solves the old IPv4 address assignment problem:**

  How can I get global IPv4 multicast addresses (GLOB, ..)

Agenda

- Why IPv6?
- IPv6 multicast addresses
- IPv6 multicast features
- Configuration

IPv6 Multicast Features

- **Group Concept**

  Multicast is based on the concept of a group.

  A multicast group is an arbitrary group of receivers that expresses an interest in receiving a particular data stream.

  This group has no physical or geographical boundaries—the receivers can be located anywhere on the Internet or in a private network.

  Receivers that are interested in receiving data flowing to a particular group must join the group by signalling their local router.

  This signalling is achieved with MLD protocol, which is the IPv6 equivalent of the IGMP protocol on IPv4.

  The network then delivers data to potentially unlimited receivers, using only one copy of the multicast data per subnet.

IPv6 Multicast Features

- **Protocol Independent Multicast v2 (PIMv2)**

  Provides intradomain multicast forwarding for all underlying unicast routing protocols

  Independent from any underlying unicast protocol such as OSPF or MP-BGP

  Sparse mode: relies upon an explicit joining method before attempting to send multicast data to receivers of a multicast group

- **Multicast Listener Discovery (MLD) v1 & v2**

  Protocol used by IPv6 hosts to communicate multicast group membership states to local multicast routers

  Version 2 of MLD adds source awareness to the protocol. This allows the inclusion or exclusion of sources.

  MLDv2 is required for Source Specific Multicast (SSM)

- **PIM Source Specific Multicast**

  SSM forwarding uses only source-based forwarding trees.

  SSM range is defined for inter domain use.

6net

GRNET
ΕΔΕΤ

IPv6 Multicast Features

- **Multiprotocol Border Gateway Protocol**

  Multiprotocol extensions to the BGP unicast inter-domain protocol that carry multicast specific routing information.

  Adds capabilities to BGP to enable multicast routing policy throughout the Internet and connect multicast topologies between BGP autonomous systems.

  Carries multiple instances of routes, for unicast and multicast routing.

- **Boot Strap Router (BSR)**

  BSR is a mechanism where a PIM router learns the set of group-to-RP mappings required for PIM SM

- **Static Rendezvous Point**

  Allow the manual configuration of the IPv6 PIM SM RP address

- **Embedded Rendezvous Point**

  Utilizes unicast based prefix addressing to include within the group address (the Rendezvous Point address)

6net GRNET ΕΔΕΤ

Agenda

- Why IPv6?
- IPv6 multicast addresses
- IPv6 multicast features
  - IPv6 Multicast Service Models
  - IPv6 Multicast & Layer2
  - Host to router signaling
  - Multicast domains
  - Protocol Independent Multicast
  - RP Operations
- Configuration

6net GRNET EAET

IPv6 Multicast Service Models

- **ASM – Any Source Multicast**

  (Traditionally just called IP Multicast)

  Service description: RFC1112 (no update for IPv6 done yet)

  MLDv1 RFC2710 **or MLDv2, PIM-Sparse Mode (PIM-SM),**

  Bidirectional PIM (PIM-bidir)

  Use ASM for legacy, dynamic- or many-source multi-party application, try to limit their use to Intradomain:

- **SSM – Source Specific Multicast**

  Service description (IPv4/IPv6): draft-ietf-ssm-overview-xx.txt

  MLDv2 required

  PIM-SSM – not a separate protocol, just a subset of PIM-SM !

  Unicast prefix based multicast addresses ff30::/12

  SSM range is ff3X::/32, current allocation is from ff3X::/96

  Use SSM for media-broadcast or interdomain applications due to simplicity and protection from DoS attacks.

6net

GRNET EAET

Agenda

- Why IPv6?
- IPv6 multicast addresses
- IPv6 multicast features
  - IPv6 Multicast Service Models
  - IPv6 Multicast & Layer2
  - Host to router signaling
  - Multicast domains
  - Protocol Independent Multicast
  - RP Operations
- Configuration

6net GRNET EAET

IPv6 Multicast & Layer2

- **Media Access Control (MAC) layer addresses within Ethernet are 48 bit addresses. These 48 bits comprise:**
  - 24 bits for the Organizational Unit Identifier (OUI) and
  - 24 bits for serial number of the card, which becomes the remainder of the unique address.

  The address of a multicast group does not relate to a physical device, but rather to a transient group of devices;therefore, the MAC address format uses a special OUI.
- **The OUI for IPv4 Multicast is 01:00:5E.** 24 bits are available for the group address – possible address overlap at Layer2**.**
- **There is a new OUI format for IPv6 Multicast:** The leading two Bytes are set to **33-33**, while the following 4 bytes/32bits are available for address mapping from the last 32 bits of the 128 bit Multicast address.

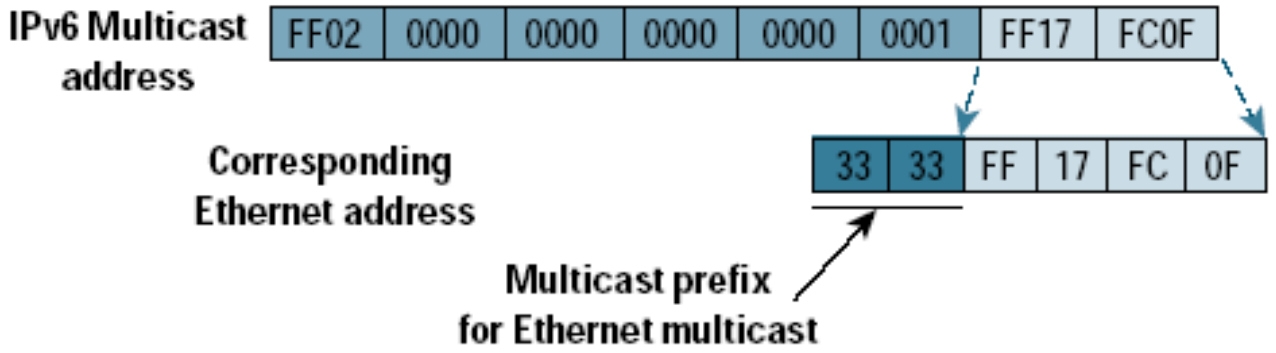

Agenda

- Why IPv6?
- IPv6 multicast addresses
- IPv6 multicast features
  - IPv6 Multicast Service Models
  - IPv6 Multicast & Layer2
  - Host to router signaling
  - Multicast domains
  - Protocol Independent Multicast
  - RP Operations
- Configuration

6net GRNET EAET

Host to router signaling

- **MLD is equivalent to IGMP in IPv4**
- **Sub protocol of ICMP:** MLD messages are transported over ICMPv6
- **MLD uses link local source addresses** (hop limit 1, router alert option)
- **Version number confusion:**

  MLDv1 (RFC2710) like IGMPv2 (RFC2236)

  MLDv2 (draft) like IGMPv3 (RFC3376)

  MLDv2 enables IPv6 to use SSM operation
- **Service Model requirements:**

  ASM – MLDv1 sufficient

  SSM – Requires MLDv2 (*Fully backward compatible with MLDv1 on hosts)*

6net GRNET EAET

Agenda

- Why IPv6?
- IPv6 multicast addresses
- IPv6 multicast features
  - IPv6 Multicast Service Models
  - IPv6 Multicast & Layer2
  - Host to router signaling
  - Multicast domains
  - Protocol Independent Multicast
  - RP Operations
- Configuration

6net GRNET EAET

Multicast domains

- A PIM domain is topology served by common RP for all sources and receivers of same group.
- A routing domain is consistent with AS.

  Its necessary to constrain the PIM messages, rp-mappings and data for groups within the PIM domain:

  – In IPv4 we used multicast boundary/ BSR border

  – In IPv6 we use scopes and zones

- SSM

  No RP or shared tree procedures (SPT only)

  MLDv2 (IPv6) required

  FF3x::/96

- ASM

  PIM-SM need RP

6net

GRNET ΕΔΕΤ

Agenda

- Why IPv6?
- IPv6 multicast addresses
- IPv6 multicast features
  - IPv6 Multicast Service Models
  - IPv6 Multicast & Layer2
  - Host to router signaling
  - Multicast domains
  - Protocol Independent Multicast
  - RP Operations
- Configuration

6net GRNET ΕΔΕΤ

Protocol Independent Multicast

- Sparse-Mode

  PIM-SM uses a pull model to deliver multicast traffic. Only network segments with active receivers, which explicitly request the data, will receive the traffic.

  PIM-SM distributes information about active sources by forwarding data packets on the shared tree.

  PIM-SM initially uses shared-trees so it requires the use of a rendezvous point

  Sources register with the RP and subsequently forward data down the shared tree to the receivers

  PIM-SM has an inter-domain deployment problem as there is no MSDP like protocol. Static use is acceptable in the intra-domain.

- SSM

  PIM SSM bypasses the shared Tree and immediately creates the Shortest Path Tree, as MLDv2 can specify the source in its request to the local router

  PIM SSM deployment may work for the inter-domain

Agenda

- Why IPv6?
- IPv6 multicast addresses
- IPv6 multicast features
  - IPv6 Multicast Service Models
  - IPv6 Multicast & Layer2
  - Host to router signaling
  - Multicast domains
  - Protocol Independent Multicast
  - RP Operations
- Configuration

6net GRNET EAET

RP Operations

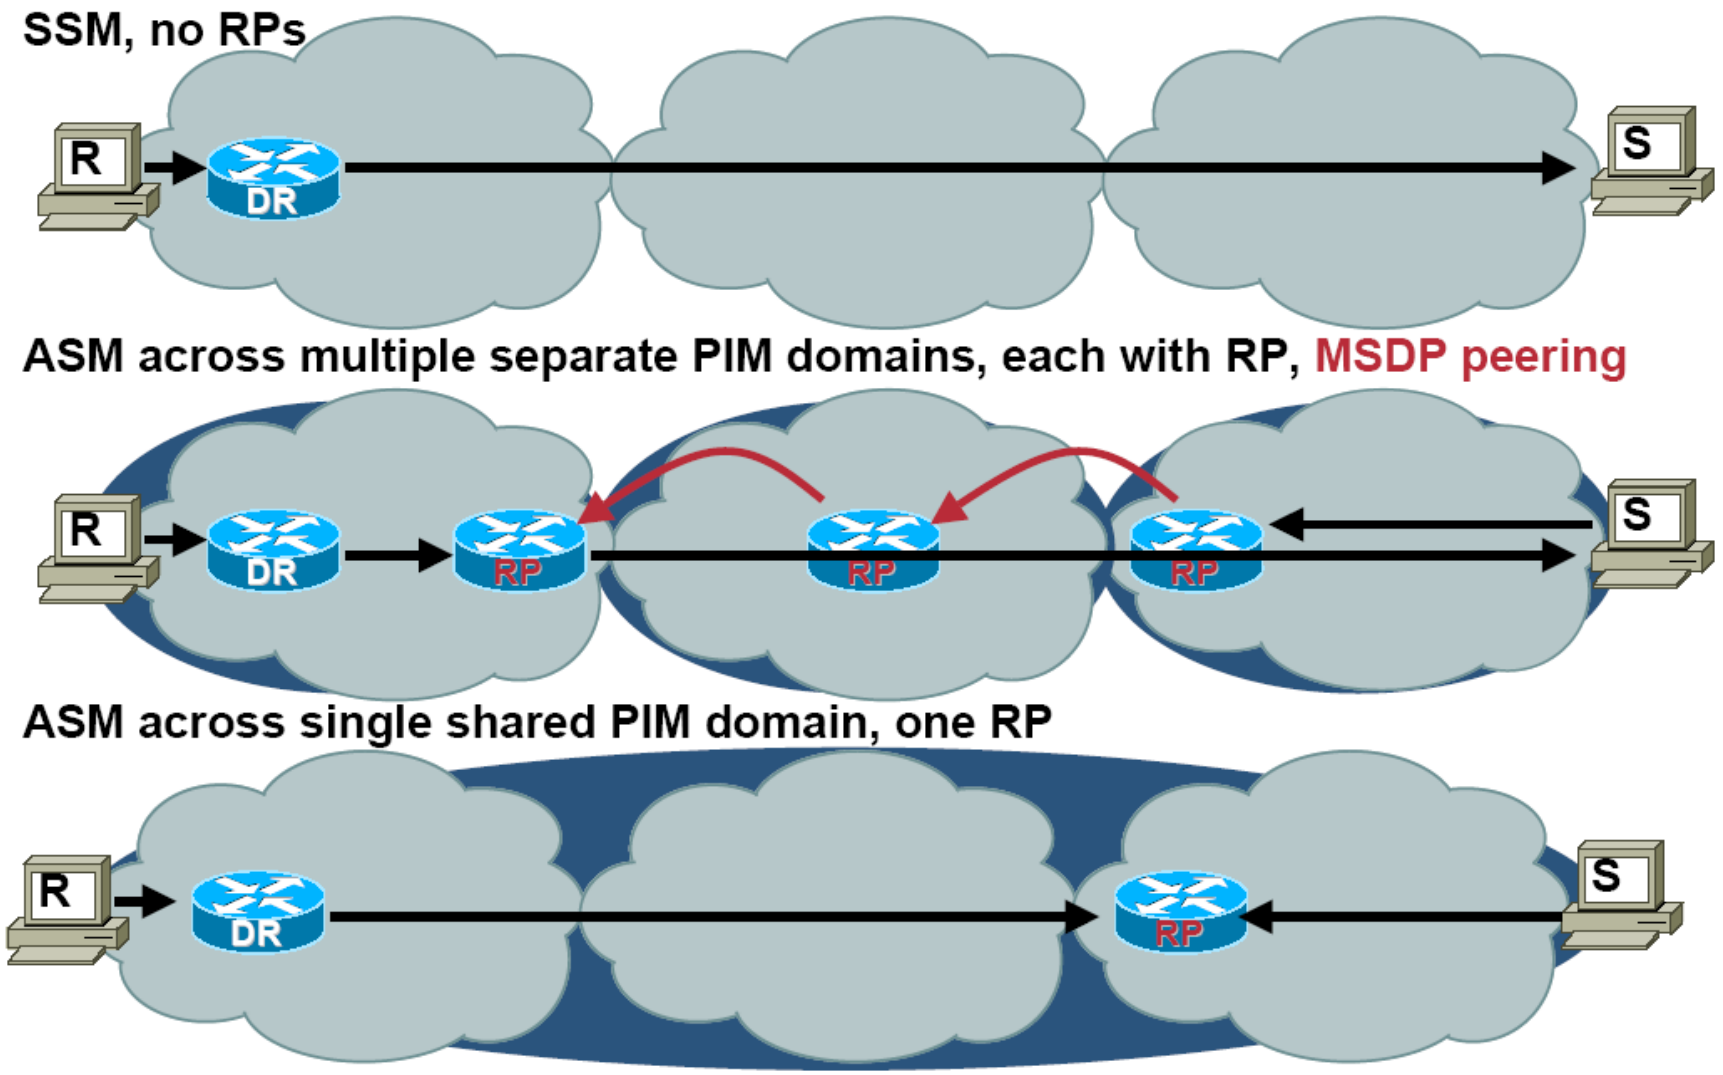

RP Operations

- **Static RP**

  For PIM-SM

  Provides Group-to-RP mapping, no RP-redundancy

  - **Easier than before as PIM is auto-enabled on every interface**

```
ipv6 multicast-routing
!
interface Loopback0
 description IPV6 IPmc RP
 no ip address
 ipv6 address 3FFE:C15:C003:110A::1/64
!
ipv6 pim rp-address 3FFE:C15:C003:110A::1/64
```

```
ipv6 multicast-routing
!
ipv6 pim rp-address 3FFE:C15:C003:110A::1/64
```

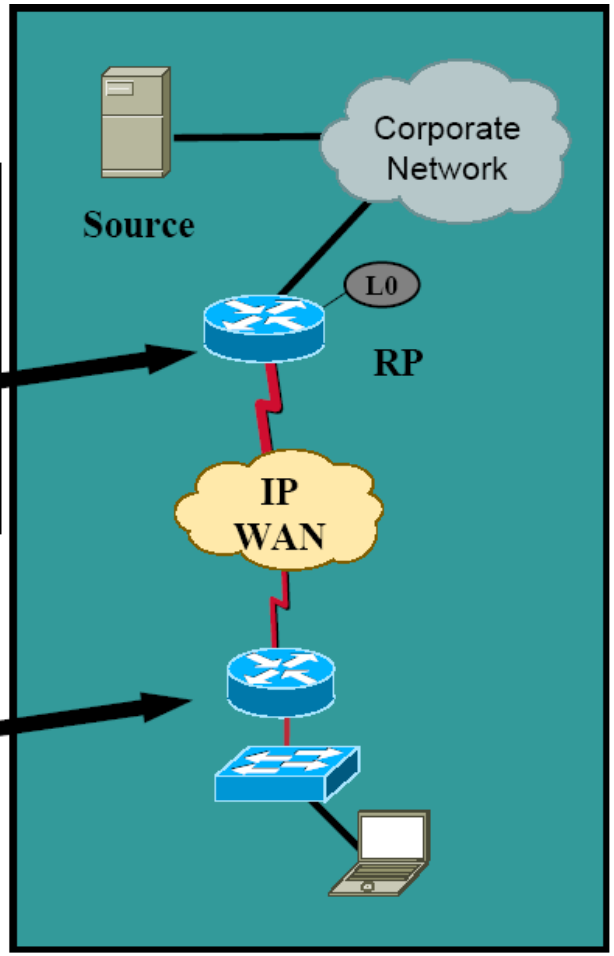

RP Operations

- **Boot Strap Router (BSR)**

  Provides Group-to-RP mapping AND RP Redundancy

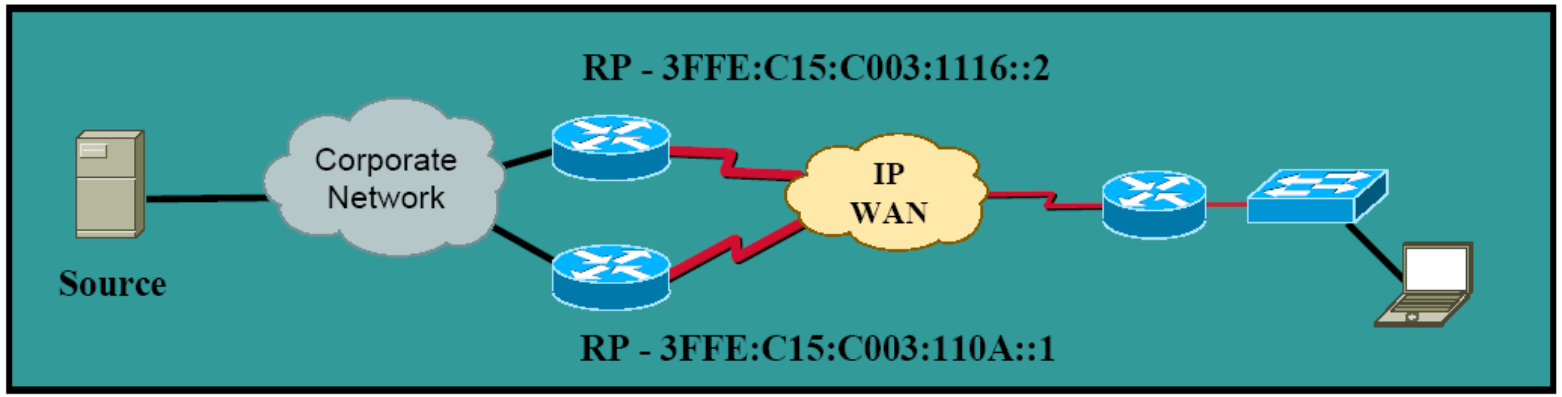


```
wan-agg-left#show ipv6 pim bsr election
PIMv2 BSR information

BSR Election Information
 Scope Range List: ff00::/8
    BSR Address: 3FFE:C15:C003:1116::2
    Uptime: 2d21h, BSR Priority: 0, Hash mask length: 126
    RPF: FE80::201:42FF:FE2D:9580,Serial2/1/0.2
    BS Timer: 00:01:44
 This system is candidate BSR
     Candidate BSR address: 3FFE:C15:C003:110A::1, priority: 0, hash mask
length:126
```

RP Operations

```
branch#show ipv6 pim tunnel
Tunnel1*
  Type   : PIM Encap
  RP     : 3FFE:C15:C003:1116::2
  Source: 3FFE:C15:C003:111E::2
```

```
branch#show interface tunnel 1
Tunnel1 is up, line protocol is up
  Hardware is Tunnel
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 3FFE:C15:C003:111E::2 (Serial0/2),
destination 3FFE:C15:C003:1116::2
  Tunnel protocol/transport PIM/IPv6, key disabled,
sequencing disabled
  Checksumming of packets disabled
  Tunnel is transmit only
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
... output truncated...
```

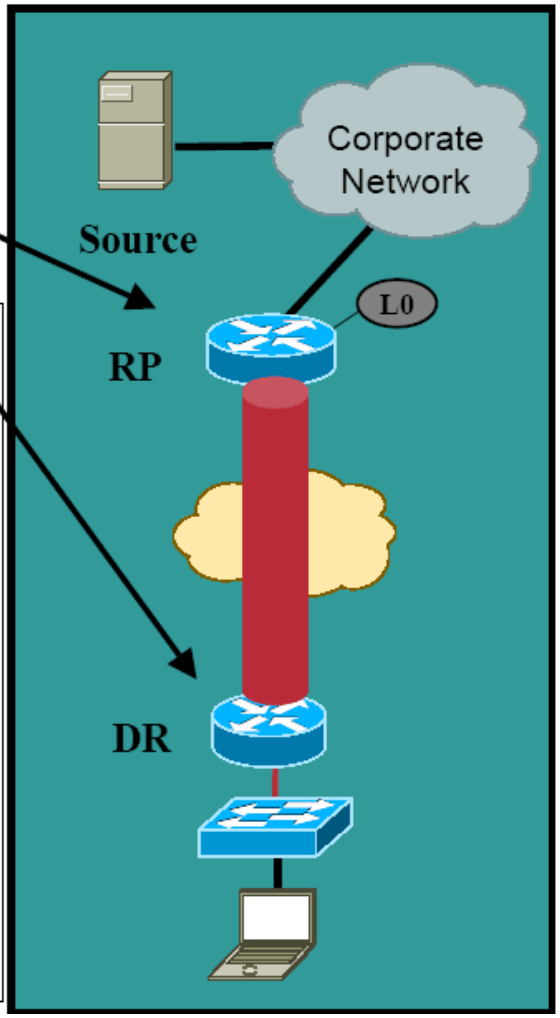

Agenda

- Why IPv6?
- IPv6 multicast addresses
- IPv6 multicast features
- Configuration

Configuration: MLD


```
6net.auth.gr#sh ipv6 mld interface fa2/0.112
FastEthernet2/0.112 is up, line protocol is up
  Internet address is FE80::208:A4FF:FEA7:7C38/10
  MLD is enabled on interface
  Current MLD version is 2
  MLD query interval is 125 seconds
  MLD querier timeout is 255 seconds
  MLD max query response time is 10 seconds
  Last member query response interval is 1 seconds
  MLD activity: 660 joins, 646 leaves
  MLD querying router is FE80::208:A4FF:FEA7:7C38 (this system)

6net.auth.gr#sh ipv6 mld traffic
MLD Traffic Counters
Elapsed time since counters cleared: 7w0d

                                  Received      Sent
Valid IGMP Packets                  2327436      458506
Queries                             55919        201374
Reports                             1107555      257136
Leaves                              244          0
Mtrace packets                      0            0

Errors:
Malformed Packets                                0
Bad Checksums                                    0
Martian source                                   190
Packets Received on MLD-disabled Interface       6
```

Configuration: MP-BGP

```
router bgp 5470
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 2001:648:0:1006::1 remote-as 5408
 neighbor 2001:648:0:1006::1 description GRNET-6NET-ROUTER
 !
 address-family ipv6
 neighbor 2001:648:0:1006::1 activate
 neighbor 2001:648:0:1006::1 soft-reconfiguration inbound
 network 2001:648:202::/48
 exit-address-family
 !
 address-family ipv6 multicast
 neighbor 2001:648:0:1006::1 activate
 neighbor 2001:648:0:1006::1 soft-reconfiguration inbound
 network 2001:648:202::/48
 exit-address-family
```

Configuration: MP-BGP

```
6net.auth.gr#sh bgp ipv6 multicast sum
BGP router identifier 155.207.199.119, local AS number 5470
BGP table version is 326611, main routing table version 326611

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2001:648:0:1006::1
                4  5408  446896  133790   326611    0    0 4d07h           71
```

```
BGP neighbor is 2001:648:0:1006::1,  remote AS 5408, external link
 Description: GRNET-6NET-ROUTER
  BGP version 4, remote router ID 2.2.2.2
  BGP state = Established, up for 4d07h
  Last read 00:00:16, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv6 Unicast: advertised and received
    Address family IPv6 Multicast: advertised and received
```

```
     Network          Next Hop            Metric LocPrf Weight Path
                                                         0 5408 6680 1717 ?
  *> 3FFE:400F:F000::/48
                      2001:648:0:1006::1
                                                         0 5408 6680 1717 2549 ?
  *> 3FFE:4013:22BB::/48
                      2001:648:0:1006::1
                                                         0 5408 6680 1717 2549 ?
  *> 3FFE:82F0::/28   2001:648:0:1006::1
                                                         0 5408 6680 1717 2549 i
  *> 3FFE:82F0:4000::/48
                      2001:648:0:1006::1
                                                         0 5408 6680 1717 2549 ?
```

Configuration: PIM-SM & BSR


```
6net.auth.gr#sh ipv6 pim neighbor
Neighbor Address            Interface          Uptime    Expires DR pri Bidir

FE80::C2B1:D2DA             Tunnel900          4d07h     00:01:40 1 (DR) B

6net.auth.gr#sh ipv6 pim bsr election
PIMv2 BSR information

BSR Election Information
  Scope Range List: ff00::/8
     BSR Address: 2001:660:3007:300:1::
     Uptime: 4d07h, BSR Priority: 255, Hash mask length: 126
     RPF: FE80::C2B1:D2DA,Tunnel900
     BS Timer: 00:02:02

BSR Election Information for Unassigned(B) Zone
  Scope Range List: ffXB::/16
     BSR Address: 2001:798:28::1
     Uptime: 4d07h, BSR Priority: 5, Hash mask length: 126
     RPF: FE80::C2B1:D2DA,Tunnel900
     BS Timer: 00:01:37
```

6net

GRNET EAET

Configuration: PIM-SM & BSR


```
6net.auth.gr#sh ipv6 pim group-map

FF3E:30:2001:700:1:FFFF::/96*
    SM, RP: 2001:700:E000:501::2
    RPF: Tu900,FE80::C2B1:D2DA
    Info source: BSR From: 2001:660:3007:300:1::(00:01:40), Priority: 4
    Uptime: 4d07h, Groups: 0
FF3E:20:2001:660::/64*
    SM, RP: 2001:660:3007:300:1::
    RPF: Tu900,FE80::C2B1:D2DA
    Info source: BSR From: 2001:660:3007:300:1::(00:01:40), Priority: 5
    Uptime: 4d07h, Groups: 0
FF3E:20:2001:660::/64*
    SM, RP: 2001:660:3007:300:1::
    RPF: Tu900,FE80::C2B1:D2DA
    Info source: BSR From: 2001:798:28::1(00:02:15), Priority: 5
    Uptime: 4d07h, Groups: 0
```

6net

GRNET
ΕΔΕΤ

Configuration: PIM-SM & BSR


```
FF3F::/32*
    SSM
    Info source: Static
    Uptime: 3w3d, Groups: 0
FF02::/16*
    L-Local
    Info source: Default
    Uptime: 3w3d, Groups: 23
FF0B::/16*
    SM, RP: 2001:610:14:5145::145
    RPF: Tu900,FE80::C2B1:D2DA
    Info source: BSR From: 2001:798:28::1(00:02:06), Priority: 100
    Uptime: 15:06:24, Groups: 0
FF0B::/16*
    SM, RP: 2001:610:14:5145::145
    RPF: Tu900,FE80::C2B1:D2DA
    Info source: BSR From: 2001:660:3007:300:1::(00:01:30), Priority: 100
    Uptime: 4d07h, Groups: 0
FF0B::/16
    SM, RP: 2001:620:0:C000::1
    RPF: Tu900,FE80::C2B1:D2DA
    Info source: BSR From: 2001:798:28::1(00:02:06), Priority: 192
    Uptime: 4d07h, Groups: 0
```

6net

GRNET EAET

Configuration: PIM-SM tunnels


```
6net.auth.gr#  sh  ipv6 pim tunnel
Tunnel5*
  Type   : PIM Encap
  RP     : 2001:610:14:5145::145
  Source: 2001:648:0:1006::2
Tunnel1*
  Type   : PIM Encap
  RP     : 2001:620:0:C000::1
  Source: 2001:648:0:1006::2
Tunnel3*
  Type   : PIM Encap
  RP     : 2001:660:3007:300:1::
  Source: 2001:648:0:1006::2
Tunnel2*
  Type   : PIM Encap
  RP     : 2001:760:600::1
  Source: 2001:648:0:1006::2
Tunnel4*
  Type   : PIM Encap
  RP     : 2001:700:E000:501::2
  Source: 2001:648:0:1006::2
```

```
6net.auth.gr#sh int tunnel 2
Tunnel2 is up, line protocol is up
  Hardware is Tunnel
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 2001:648:0:1006::2 (Tunnel900), destination 2001:760:600::1
  Tunnel protocol/transport PIM/IPv6, key disabled, sequencing disabled
  Checksumming of packets disabled
  Tunnel is transmit only
```

Questions

6net

GRNET
ΕΔΕΤ